

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
15 April 2004 (15.04.2004)

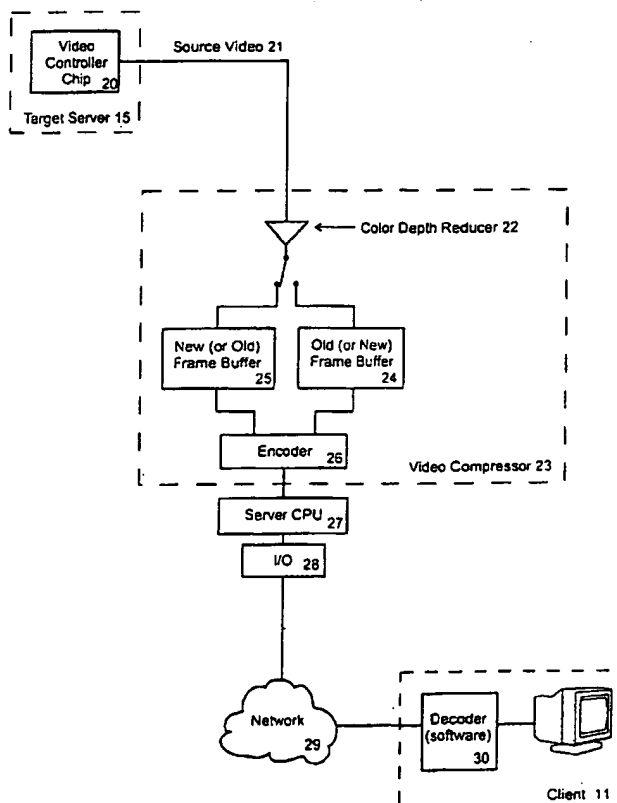
PCT

(10) International Publication Number
WO 2004/032356 A1

- (51) International Patent Classification⁷: **H04B 1/66** (81) Designated States (*national*): AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, UA, UG, UZ, VN, YU, ZA, ZW.
- (21) International Application Number: PCT/US2003/010488
- (22) International Filing Date: 7 April 2003 (07.04.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 10/260,534 1 October 2002 (01.10.2002) US
- (71) Applicant: **AVOCENT CORPORATION** [US/US]; 4991 Corporate Drive, Huntsville, AL 35805-6201 (US).
- (72) Inventor: **DAMBRACKAS, William, A.**; 3040 Birkdale, Weston, FL 33332-1810 (US).
- (74) Agent: **DAVIDSON, Scott, J.**; Davidson Berquist Klima & Jackson LLP, 4501 North Fairfax Drive, Suite 920, Arlington, VA 22203 (US).
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published: — with international search report

[Continued on next page]

(54) Title: VIDEO COMPRESSION SYSTEM



(57) Abstract: A video compression system (Fig. 2) is disclosed that is optimized to take advantage of the types of redundancies typically occurring on computer screen (11) and the types of video loss acceptable to real time interactive computer users (11). It automatically adapts to a wide variety of changing network (29) bandwidth conditions and can accommodate any video resolution and an unlimited number of colors. The video compression encoder can be implemented with either hardware or software and it compresses the source video into a series of data packets that are a fixed length of 8 bits or more.

WO 2004/032356 A1



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

TITLE OF THE INVENTION
VIDEO COMPRESSION SYSTEM

FIELD OF THE INVENTION

[0001] This invention relates to computer data processing, and more particularly to computer video compression.

BACKGROUND OF THE INVENTION

[0002] Existing video compression systems can compress a stream of video data so it takes less bandwidth to send over a communication channel. Such systems take advantage of redundancies expected to occur in the video they are aiming to compress. For example, JPEG and MPEG take advantage of frequent similarities in the colors of adjacent pixels in photographic images. In addition, MPEG takes advantage of the fact that motion pictures often have many pixels that stay the same color during many frames of video or only shift their positions on the screen as the camera moves.

[0003] Video can be further compressed depending on how much degradation in video quality (or "video loss") is acceptable to the person (or "user") viewing the video, but the acceptability of different types of video loss is highly dependent on the user's activity (or "application"). The four types of video loss are; (1) resolution loss (appears blurred), (2) color depth loss (has fewer shades of colors), (3) frame rate loss (stalling or jerkiness of a motion picture) and (4) time loss or "video delay" (time delay from video capture to its availability for viewing).

[0004] To achieve higher compression ratios, different compression systems take advantage of the types of video loss that are the most acceptable to the users they aim to satisfy. For example, with MPEG, fast action scenes that would generate too much data

for the communication channel are sent with resolution loss because movie viewers accept resolution loss better than they accept frame rate loss or color depth loss.

[0005] Video delay is not a problem in some applications but it is a serious problem in other applications. Different compression systems impose different amounts of delay as they compress the video. Systems that impose more delay achieve higher compression ratios because all the video frames captured, held and examined during the delay provide a better opportunity to decide how to compress them. One example might be: "is the camera moving or is just one object in the scene moving".

[0006] Video delay is not a problem with "one-way" user activities, such as watching movies; therefore, the compression systems used for these applications (such as MPEG) impose a long delay (many seconds or more) before compressing the video and beginning to send it over the communication channel. In fact, when the communication channel is a network with indeterminate bandwidth availability (such as the Internet), the video received from the network is often buffered and delayed for many more seconds before it is displayed (to eliminate the stalling caused by network congestion). Although time delay is not a problem with one-way user activities such as watching movies, it is a serious problem for real time "interactive" users, such as users with a mouse, controlling a cursor that is a part of the compressed video image.

[0007] One such example of real time interactive users relates to the remoting of a computer KVM console (Keyboard, Video display and Mouse) over a communication channel. In these "remote console" applications, keyboard and mouse data are sent from the remote console over the communication channel and "switched" to one of a number of "target" server computers, just as if the keyboard and mouse were directly connected to that target server. The corresponding video is sent from the target server to the remote console just as if the target server was directly connected to the remote console's video display. Examples of KVM systems are described in commonly-owned U.S. Patents Nos. 5,721,842 to Beasley et al and 5,732,212 to Perholtz et al, each of which is incorporated herein by reference.

[0008] The communication channel for some KVM systems provides enough bandwidth to transport the uncompressed video because they use dedicated local cables

and a dedicated circuit switch. KVM systems adapted to operate over a network via, for example, Internet protocol (referred to herein for brevity as "KVM/IP" systems) provide limited and indeterminate bandwidth availability compared to a dedicated local cable-based KVM system. Sending keyboard and mouse information from the remote console to the selected target server in a timely fashion is one concern with KVM/IP systems. A greater concern is sending the relatively high volume of video data back to the remote console in a timely fashion. Since today's typical computers output video continuously at over 2 gigabits per second and remote Internet connections (such as DSL) typically operate at less than 1 megabit per second, video compression ratios averaging well over 2000-to-1 are required. Remote Internet connections using dial modems at 50 kilobits per second require even higher average compression ratios.

[0009] As a remote console user moves their mouse or types on their keyboard to input new information to the server, those actions must be communicated to the server and acted upon by the server to create new video images, which are sent back to the remote console user's screen. Delays in sending the video back to the remote console user are annoying because they create a temporal lag between the entry of the keyboard or mouse information by the user and the video response perceived by the user on their screen. Delays following keyboard activity are less annoying than delays following mouse movements, thus the term "mouse-cursor response" is used to describe this problem.

[0010] This problem of remote console applications (described above) is not applicable to some types of typical web browser applications. With web browser applications, the video cursor image is created locally on the user's computer, so mouse-cursor response is always very good even if the network is slow at responding with server-generated video images. With remote console applications, network delays affect the mouse-cursor response because the cursor is represented as an integral part of the video image coming from the server and sent to the remote console over the network.

[0011] In remote console applications, user acceptability for the four types of video loss is the complete opposite from other video applications. As described above, minimum video time delay is a factor in remote console applications, but video delay is a

less important type of video loss in other applications. The importance of resolution loss in remote console applications is also the opposite of other applications because the computer screens sent to remote consoles are typically made up of a significant amount of relatively small font alphanumeric text, many small icons and many high contrast sharp edges. Compression systems such as JPEG or MPEG, that impose resolution loss may be satisfactory for many other applications, but they are not satisfactory for reading small font alphanumeric text and images with high contrast sharp edges. The opposite order of user acceptability also applies to color depth loss and frame rate loss. These two types of video loss are the most acceptable by users in remote console applications and the least acceptable in other video applications.

[0012] Although existing video compression systems are widely used and well suited for a wide variety of applications, a video compression system optimized for the best possible interactive computer user experience is needed.

BRIEF SUMMARY OF THE INVENTION

[0013] The present invention is a new video compression system that is optimized to take advantage of redundancies typically occurring on computer screens and also is optimized to take advantage of types of video loss acceptable to real time interactive computer users. In one embodiment of the present invention, captured frames of computer video are "encoded" into combinations of five different, uniquely chosen "commands", which are selected and sequenced based on their ability to most efficiently compress the captured video. These commands are sent over the network to the "client" where they continuously instruct (or command) the "decoder" on how to decompress or decode the commands and recreate the captured video frames on the remote video display. In a unique way, this embodiment can compress and decompress computer video without resolution loss or color depth loss, but with frame rate loss that is dynamically adjusted depending on available network bandwidth. It also imposes minimal delay during encoding and decoding.

[0014] The five commands are; (1) copy old pixels from an earlier frame (sometimes called "no change from an earlier frame," "no change" or simply "NC"), (2) copy pixel from the left, (3) copy pixels from above, (4) make a series of pixels using a 2-color set, and (5) make one or more pixels using a specified color. Each command provides unique efficiencies when employed in a hierarchical structure. Also, the commands are included in comprised of packets that are a fixed length of 8 bits or more, such that they can be easily sent, received and decoded with either software or hardware. The present invention is not limited to any command or packet length, but preferred embodiments would use lengths that were a multiple of 8-bits (such as 16, 32 or 64) such that they would be compatible with popular and commonly available components and processors.

[0015] In broader embodiments of the present invention, one, two, three, or four of the types of commands described above are used, alone or in any combination thereof. For example, the inventor believes that the use of the command to make a series of pixels from a 2-color set alone is unique in compressing video that includes a significant amount of alphanumeric text (such as viewing this document with a word processing program). Further advantages and efficiencies are gained when others of the commands are added thereto in various combinations. In other embodiments, one, two, three, four, or all five of the commands are used in conjunction with any kind of prior art compression system to enhance the video compression of the known system. For example, MPEG, JPEG, and others (and all variants thereof (e.g., MPEG2, etc.)) can be used with one or more of the five commands described herein to enhance the video compression of the prior art compression techniques.

[0016] In other embodiments of the invention, referred to as the "gray-favored" color modes, the captured video can be further compressed by taking advantage of the fact that remote console users accept color depth loss better than any other type of video loss. In this mode, each pixel of the captured video is converted to the nearest color from a set of specifically chosen colors that match typical colors used on computer screens. Grays are favored in the set of colors since they are favored on typical computer screens (white and black are included in the definition of "grays").

[0017] The present invention can be embodied with the compression encoding implemented with hardware, with software or with a combination of hardware and software. Likewise, the decoding can be implemented with hardware, with software or with a combination. The "source" video can be captured by connecting directly to a video controller chip inside a computer. Alternatively, the video can be captured from a computer's external analog video output, external digital video interface (DVI) or other external interface. In one embodiment, the video is compressed with hardware using an FPGA (field programmable gate array) or ASIC (application specific integrated circuit). In another embodiment, the video is compressed completely with software before it is made into a video output stream.

[0018] The video compression commands are sent over a network to the remote console where they are decompressed and displayed to the user. The remote console can be a conventional PC (personal computer), which decodes the commands using PC software or it could be a small "thin client" device built with a low performance microprocessor. In one embodiment, the commands are all designed to consist of one or more 8-bit packets so that they can be easily decompressed with software running on a low performance microprocessor. Alternatively, a hardware device (such as an FPGA or ASIC) can completely decode the commands at the remote console. In such a case, the remote console would not require a computing device for command decoding or a video controller chip for displaying the user's video. Such a low-cost hardware (or combined hardware/software) remote console is referred to below as a "microclient."

[0019] The present invention also has application in computer "blade" technologies, where individual server computers are contained on single cards, and many of these cards are assembled into a common blade chassis to share a common power supply and central control functions. Conventional cable-based KVM switch technology on the blades can give local cable-attached users access to each blade computer, but if users need KVM access to the blades over a network, the present invention could be included in the blade chassis or on each blade and the video compression commands could be given to a common network interface in the blade chassis to be sent over the network to the various remote consoles.

[0020] This invention can thus be employed in generally compressing computer video, for sending computer video over LANs, WANs, dial-up or any other networks, for applications in thin client, microclient, and remote console applications (such as KVM/IP systems).

BRIEF DESCRIPTION OF THE DRAWINGS

[0021] The patent application file contains at least one drawing executed in color. Copies of this patent or patent application publication with color drawings will be provided by the Office upon request and payment of the necessary fee.

[0022] FIGURE 1 is a schematic representation of an example embodiment of the present invention in a KVM/IP system with the client implemented with PC software;

[0023] FIGURE 2 is a schematic representation of an example embodiment of the present invention showing the internal operation of a hardware video compressor ;

[0024] FIGURES 3-10 are tables of the video compression commands in an example embodiment of the present invention with 8-bit packet lengths;

[0025] FIGURE 11 is a chart describing how color depth is reduced in the "7-bit gray-favored color mode" embodiment of the present invention;

[0026] FIGURE 12 is a color print of test pattern (referred to as the 0-255 RGB+Gray test pattern) on the video screen of a computer set for 24-bit color.

[0027] FIGURE 13 is a color print of a client computer screen when the "7-bit gray-favored color mode" embodiment of the present invention is employed and the source video is the test pattern shown in Figure 12;

[0028] FIGURE 14 is a schematic representation of an example embodiment of the present invention with the video creation software and the video controller chip integrated together with the video compressor;

[0029] FIGURE 15 is a schematic representation of an example embodiment of the present invention without a video controller chip, and with software video compression;

[0030] FIGURE 16 is a schematic representation of an example embodiment of the present invention referred to as a microclient;

[0031] FIGURE 17 is a schematic representation of an example embodiment of the present invention describing the concept of "share-mode";

[0032] FIGURE 18 is a chart describing how color depth is reduced in the "5-bit gray-favored color mode" embodiment of the present invention; and

[0033] FIGURES 19-24 are tables of the video compression commands in an alternative embodiment of the present invention for use with the 5-bit and 12-bit color modes.

DETAILED DESCRIPTION OF THE INVENTION

[0034] The present invention can be implemented with any hardware or software that aims to send computer video over a communication channel, including over an intervening network. One such example embodiment is shown in Figure 1, which is described by way of example rather than limitation. Indeed other embodiments will be understood once the artisan reviews the invention, as embodied in the appended Figures and described herein.

[0035] In Figure 1, the KVM/IP system 10 includes a remote console client 11 and a server appliance 14. In the embodiment shown, the remote console 11 is implemented with PC software in a network-ready PC (which includes a keyboard, video display and mouse). The client 11 communicates via an Internet Protocol network 13, to a target server 15 via the KVM/IP appliance 14. The appliance 14 and the client 11 include standard network I/O hardware and software to permit them to communicate via any form of Internet protocol connectivity, such as a dial-in, DSL, WAN, LAN, T1, wireless, etc.

[0036] In Figure 1, the appliance acts as an intermediary between the target server 15 and the client 11, permitting the client 11 to couple its keyboard, video display and mouse to the server 15 just as though the client 11 was directed connected to it. In that regard, the manner and operation of the system 10 combined with the addressing and switching capability of the IP network is typical of KVM switches such as those sold by the present assignee, and by Cybex Computer Products of Huntsville, Alabama, and by Apex, Inc. of Redmond, Washington.

[0037] The client 11 includes software that facilitates the identification of the target server 15 (via the appliance 14) such as by a standard TCP/IP address. Once communication is established between the client 11 and the appliance 14, the client 11 employs software to send keyboard and mouse data, entered at the client, to appliance 14 via the IP network 13. The appliance 14 receives the data switched or routed to it, and applies it to the keyboard and mouse ports of the server 15 just as if the keyboard and mouse were directly attached to the server 15. In response, the server 15 acts on the keyboard and mouse data (via whatever application is running on the server 15) to produce new video data, which is output to the appliance 14 via the video output of the server 15.

[0038] Once the appliance 14 receives the video from the server 15, it compresses it by one of the inventive algorithms described below and transmits the resulting video compression commands to the client 11 via IP network 13. Compression can be done with an FPGA, ASIC, or any other hardware or software in the appliance 14. Alternatively, appliance 14 can be "embedded" into the server 15, or it can be eliminated if server 15 includes software to perform the compression and send the resulting commands directly to the IP network 13. Upon receipt, the client 11 decodes the commands with PC software and reproduces the target server's video on the client PC's screen for viewing by the user. Alternatively, the command decoding could be done with hardware in the client 11.

[0039] In the embodiment of Figure 1, the user should perceive the client PC's keyboard, video display and mouse as being directly connected to the server 15, even though the client 11 and the server 15 may be physically located as far away as opposite

ends of the globe. Imposing too much delay in getting the keyboard and mouse data to the server 15 via the appliance 14, and in getting the video back can impede that objective. The keyboard and mouse require a relatively small amount of data traffic, which can be transported quickly and relatively efficiently, but the high volume of video data presents a more difficult problem. To be effective, the video must be compressed by the appliance 14, transmitted via the IP network 13, decompressed by the client 11, and presented on the user's screen as quickly as possible. Excessive delay is most evident in mouse-cursor response. Even slight lags between mouse movements and the cursor response presented on the screen are annoying to the user.

[0040] Figure 2 illustrates an example embodiment of the present invention. There are many different hardware and software implementations that the present invention can be envisioned within and the embodiment of Figure 2 is not the only such way. After reviewing the present teaching, the artisan will recognize other ways consistent with the breadth of the present invention in which to implement the invention.

[0041] At the top of Figure 2, the source video 21 can be in any form, analog or digital. Most current video controller chips have their video output available digitally for use with flat panel displays, such as used in laptop computers. The video compressor 23 could connect directly to the output pins of the video controller chip 20 or it could connect to an external connector on the target server 15. One type of external connector is DVI (digital video interface), which is a standard for connecting digital video to external digital display devices. Any other type of source video will suffice as well—the invention is not limited as such.

[0042] Optionally, a color depth reducer 22 can be included in the video compressor 23 to reduce the number of bits defining the color of each pixel. It does this by categorizing pixels' colors into zones. When the source video 21 is digital video, the simplest means of color depth reduction is to ignore the least significant bits. For example, 24-bit color could be converted into 15-bit color by ignoring the least significant 3 bits of each of the 8-bit red, green and blue signals. Ignoring the least significant 4 bits of each 8-bit color signal would result in 12-bit color. More complex color reduction methods referred to as the 7-bit gray-favored color mode and the 5-bit

gray-favored color modes are described further below and illustrated in Figures 11 and 18.

[0043] If the source video 21 is an analog video signal, the color depth reducer 22 needs to include an A-to-D (analog to digital) converter. With analog video, each pixel is defined by three analog signals (red, green and blue). The A-to-D converter digitizes the intensity of each pixel's three signals by detecting what "zone" they are in (very similar to what the digital color depth reducer described above does). A major difference with analog video is noise. When an analog signal is on the edge of a zone, a small amount of analog noise can make the digitizer bounce back and forth from one zone to another in subsequent frames. In such a case, it will appear that the source video 21 is changing even if it is not. Consequently with an analog input, some method of noise suppression needs to be used to reduce this "zone bouncing." Any noise suppression techniques can be used, but in one example, when the input signal is within a zone, it must get out of that zone by at least a threshold amount before it is considered to be in another zone. This comparison of what zone each pixel's signals were in during the previous frame is done for every pixel in the video frame.

[0044] Although the several embodiments mentioned for the source video are contemplated within the present invention, the particular example embodiment in Figure 2 presumes digital video received from a video controller in the target server 15 will be the source video. The output of the video chip 20 is the source video 21, which is a continuing stream of video data. The video controller chip 20 need not be controlled by any aspect of the present inventions (though the inventions can certainly be employed in conjunction with some video chip control), that is, the video chip 20 will output video in a continuing stream in accordance with its own internal timing.

[0045] The source video 21 is the input to the video compressor 23. Of course, other processing devices, such as general or special purpose processors can be substituted for the hardware video compressor 23. The video compressor 23 includes at least two frame buffers 24 and 25, and may include many additional frame buffers or frame buffer types for additional operational complexities and efficiencies. Prior to the client 11 establishing a connection over the network 29, the source video 21 is continuously

captured (and continuously overwritten) in one of the frame buffers 24 or 25 (at the instant shown in Figure 2, frame buffer 25 is active, meaning it is capturing the video).

[0046] When a client initially connects over the network 29, the video capturing stops and the encoder 26 begins reading and compressing the captured video data in buffer 25. It starts at the beginning of the frame buffer (which is the upper left pixel on the screen) and progresses pixel-by-pixel to the end of the frame buffer (which is the lower right pixel on the screen), looking ahead and building the most efficient sequence of commands. As the encoder 26 builds this series of commands (in accordance with the algorithm embodiments described below), the server CPU 27 is sending them to the client 11 via the I/O 28 and the network 29. After the encoder 26 is finished with the last pixel in the buffer 25, the frame buffers switch and source video begins to be captured in the other frame buffer (buffer 24 in this case). This switch occurs even if the CPU 26 has not finished sending the commands to the network 29. After the switch, the frame in buffer 25 becomes the "old" frame and represents the frame displayed (or soon to be displayed) on the client's screen.

[0047] Since the source video was continuing to run while it was not being captured, it might be half way down the screen or anywhere else in the screen when the capturing begins. Regardless of where the new capture into buffer 24 starts, it continues for one full lap until it gets back to the screen position from which it began capturing. The result is one full "new" frame of video captured from the source video 21. If the CPU 27 has not been able to send all the commands from the first compressed frame over the network (possibly due to network congestion or a slow network) after the new frame of video is captured, then the capturing process will continue overwriting the captured video in buffer 24. When the network is ready for more commands (and at least one frame of video has been captured), the capturing will stop and the same process that occurred for the first frame will continue. However, since the client 11 now has its first frame, the encoder 26 will now be able to compare each pixel in the new frame with each pixel in the old frame and if pixels didn't change, the compression will be much better. The same process now continues after at least one frame of new video has been captured and the network is ready for more commands. This process of continuing to capture

while waiting for the network to be ready lowers the effective frame rate to the client depending on network conditions and displaying the "newest" video takes precedence to displaying "all" of the video. In effect, captured video becomes an expiring commodity. The remote console users accept frame rate loss much more than the video delay they would have to tolerate if "all" the video motion was queued and sent later.

[0048] Thus, in the present example, the new frame buffer (formerly the old frame buffer) captures the most recent frame of source video. Then, the old frame (in the old frame buffer) and the new frame (in the new frame buffer) are read by the encoder 26 for the purpose of comparing and compressing the video. There are alternative methods of capturing and comparing video frames for compression and all such methods will not be described here.

[0049] In the narrower of the embodiments of the present inventions, all aspects of the video encoding described herein with respect to Figure 3 are employed. The detailed description of all of those aspects described herein with respect to "the invention" or "the inventions" should not be construed to mean that the invention requires every aspect of the example algorithms described. The examples are provided for purposes of describing one example way in which the efficiencies of the inventions can be realized. Other, broader and narrower, aspects of the invention may be realized from the descriptions that follow. Thus, in Figure 3, five video compression commands are presented for compressing the video read from frame buffers 24 and 25. They are, in hierarchical order, (1) copy old pixels from an earlier frame, (2) copy the pixel from the left, (3) copy the pixels from above, (4) make a series of pixels using a 2-color set, and (5) make one pixel using a specified color. The inventor has discovered that this combination of hierarchical commands provides compelling video compression for computer displays. The first three of these commands provide 3 dimensional copying (horizontal, vertical and time) and the fourth command provides unique efficiencies for screen segments that are comprised of only two colors (such as text).

[0050] In the embodiment illustrated in Figure 3, there are five different video compression commands. All of the commands are made from either single packets or multiple packets, where each packet consists of one 8-bit byte. The first one to three bits

of the first packet of each command is the operation code (or "opcode") and they determine the command's basic function. The "E" bit is the "extension" bit and the rest of the bits (R, C and P) are the "payload" bits. The general formats of the five commands are shown in Figure 3, and the more detailed formats of them are shown in Figures 4-10. For embodiments with different packet lengths, the number of payload bits would be different. For example, 16-bit packets would have 8 additional payload bits.

[0051] The lowest hierarchical command, the MP (make pixel) command has a one in the first bit location (bit position seven) followed by payload bits ("P" bits) that define a color (none of the other commands begin with a one). If the number of color bits used is seven, the MP command is one byte long (as shown in Figure 3). If the number of color bits used is fifteen, the MP command will be two bytes long, with the first bit of the first byte being a one (as shown in Figure 4). Likewise, if the number of color bits (P bits) is 23, the MP command will be three bytes long (as shown in Figure 5), and so on. The MP command is the simplest command to understand, and also provides the least compression. It says, essentially "make one pixel this color" with the payload identifying the color. A popular setting for computer consoles is 15-bit color (5 bits for red, 5 for green and 5 for blue). 15-bit color would be supported by two-byte MP commands. Since single-byte MP commands have seven payload bits, they can present 2^7 (or 128) different colors. The 7-bit gray-favored color mode described further below describes how the source video can be "reduced" to the nearest of 128 colors widely used on computer consoles. The following discussion of the present invention's operation describes the operation with one-byte MP commands but the invention is not limited to a specific number of color bits (P bits).

[0052] In terms of compressibility, a frame where every pixel is a random color would be non-compressible without resolution loss (other compression systems, such as JPEG, fractal analysis, etc. could provide compression with varying degrees of resolution loss). With the embodiment of Figure 3, every single pixel in this random frame would be encoded with an MP command and if this frame had one million pixels, it would take one million MP commands to encode it. If the encoder cannot use any other command to encode a pixel, it uses an MP command. Every pixel will always qualify to be encoded

with an MP command. The MP command thus takes its place in the lowest hierarchical position in Figure 3. Being a listing of priority, Figure 3 indicates that the encoder 26 tries to make the top command, then the second, then the third, then the fourth, and then it gets to the MP command as a last resort.

[0053] Looking now at the opcodes in Figure 3, a "one" in bit-position seven uniquely identifies a make-pixel command. If a "zero" is in bit-position seven, that command is one of the other four commands shown in Figure 3, with the next two bits (bit positions five and six) identifying which of the other four commands applies. Thus, a 00 in bit locations five and six indicates a CO (copy old or no change) command, a 01 indicates a CL (copy left) command, a 10 indicates a CA (copy above) command, and a 11 indicates a MS (make-series) command. Thereafter, each of these four command types has payload bits that follow the opcode. The payload bits are the R, C and P bits. The E bits will be discussed below under the MS command.

[0054] The payload bits (R bits) in the CO, CL and CA commands indicate the number of times the command operation is repeated. The CO command instructs the client that pixels have not changed from pixels currently displayed. Thus, the encoder 26 compares the old and new frame buffers and evokes CO commands when it determines that current pixels in the "new" frame are no different from pixels at the same locations in the "old" frame. Thus, CO commands are sent for portions of the screen that are not changing in the source video.

[0055] The next two commands compare pixels in terms of locations within a common "new" frame, rather than as between the old and new frame. The CL command instructs the client to copy the color from the pixel in the position immediately to the left in the current frame. If the current pixel is the first pixel on a video line, the pixel immediately to the left is the last pixel on the previous line. The CA command instructs the client to copy the color from the pixel immediately above in the current frame. The CL, CA and CO commands are referred to below as "copy" commands. Other commands may be substituted that provide copying of pixels with relations within a common frame or as between old and new frames. The presently described commands have particular advantage in computer video because of the proliferation of horizontal

and vertical rectangles and lines that exist in computer video. With horizontal lines, for example, CL commands have particular utility and with vertical lines, CA commands have particular utility.

[0056] The final command is the MS or make-series command and is itself unique in the present types of video encoding. The MS command takes advantage of a particular aspect of computer video, namely that large portions of typical computer screens are composed of only two colors. The classic example of that in computer video is text information in which large portions of the screen are made from with a text foreground color on a solid background color. In such cases, the MS command permits the encoder 26 to create a substantial amount of the video without loss of sharpness in the text, and with very substantial amounts of compression.

[0057] Each of the commands will now be discussed in the context of their payload structures and in the context of real applications. As previously described, the CO command (Figures 3, 6 and 7) essentially identifies that a present pixel has not changed from the pixel located in the same location of the previous frame. To further the compression, the payload also identifies not only that a present pixel has not changed, but also that some number of consecutive pixels didn't change. What that number will be is described below. As shown in Figure 3 for the CO command, after the three-bit opcode, there are five bits (RRRRR) that indicate the repeat count of that CO command. These five bits can be set to any binary value between 0 and 31.

[0058] Since a repeat count of zero doesn't make sense, one would initially assume that these five bits count define up to 32 consecutive pixels in a row that did not change from the previous frame. However, if one-byte MP commands are only being used (instead of two or more byte long MP commands) a repeat count of one also does not make sense, since a one-byte make pixel (MP) command has the same compression value as an CO command with a repeat count of one. In that case the repeat count payload could start with a count of two, such that a payload of 00000 means a repeat count of two and a payload of 11111 means a repeat count of thirty-three. With that, a small additional efficiency is provided, namely that an CO command with a five-bit

payload identifies the fact that somewhere between two and thirty-three pixels have not changed from the frame displayed already.

[0059] The preferred embodiment adds still further efficiency. Suppose that more than thirty-three pixels have not changed. As shown in Figure 6, a second immediately consecutive byte with a 000 opcode can follow a first byte with 000 opcode, providing a second five bits to represent from two to thirty three pixels again. But, instead, the decoder 30 will detect two consecutive packets with CO opcodes, and interpret both the five bit payloads as a single CO command with a ten-bit payload. With a ten-bit payload, the number of consecutive CO pixels extends from 34 to 1025. In other words, with just two eight-bit bytes, a frame of over one thousand pixels can be sent to the client. The CO command is escalating in its efficiency. One can note that there is no other reason to make two consecutive packets with CO opcodes other than the fact that a repeat count over 33 is required. The encoder 26 will not make two consecutive packets with CO opcodes if a repeat count over 33 is not required.

[0060] A two-byte CO command gets inefficient briefly if the encoder 26 requires a repeat count of 35 or 36, requiring a second byte. But, once the repeat count gets up to a thousand pixels (such as a full line on a 1024x768 resolution screen), just two bytes can compress the whole line. Further, if a third CO command follows the second (as shown in Figure 7), the decoder 30 detects a fifteen-bit payload. If there's a fourth CO command, it detects a twenty-bit payload. A four-byte CO command can identify that over one million pixels have not changed, which is more than is needed for a complete frame with 1024x768 resolution. The present invention is not limited to any particular number of consecutive CO commands or any video screen resolution, but for present purposes, a five-byte command (that supports up to 33 million pixels) provides a repeat count large enough for a full frame at the highest video screen resolutions currently envisioned.

[0061] The CL and CA commands operate the same as the CO command described above. They duplicate different pixels (pixels to the left, or pixels above) but they have the same structure, a three-bit opcode followed by a 5-bit RRRRR payload identifying a repeat count. Again, each of the CL and CA commands can be sequenced,

as shown in Figure 8 for the CL command, to form 10 bit, 15 bit, 20 bit, or longer payloads.

[0062] The hierarchical priorities between the CO, CL and CA commands only apply if two or more of those commands simultaneously qualify on the current pixel. If the encoder 26 determines that the CO command qualifies on the current pixel and no other copy command qualifies, the encoder temporarily ignores the other copy commands and continues to compare pixels from the old and new frames, to determine how many pixels in a row the CO command qualifies for. The encoder 26 would do the same thing if it discovered that the CA or CL commands alone qualified on a current pixel. At the first instance that the identified (CO, CA or CL) condition is no longer true, the encoder 26 sends one or more consecutive commands of Figure 3 and then evaluates the next pixel for encoding. In other words, once the encoder 26 determines that one repeat count condition is true for a given pixel, and only one repeat count condition is true for a given pixel, it ignores all other command evaluations until the current repeat count condition is no longer valid. When that occurs, it creates the command (opcode and repeat count) and sends it to the client.

[0063] As long as one copy command (CO, CL, or CA) qualifies, the encoder continues with it until it no longer qualifies. Then the encoder ends that analysis and creates the appropriate bytes. If, however, multiple repeat count conditions (CO, CA or CL) initially qualify on the same pixel, the encoder just starts counting consecutive pixels for which those conditions apply. As long as one of these commands qualifies, the counter continues to run. Eventually, the encoder will choose only one command that applied for the full repeat count so it only counts one counter. It does not need to run three different counters, one for each copy command. Then, as the encoder continues to count, it will discover that some commands no longer qualify. When that occurs enough times so that no command type is "left standing," the encoder 26 creates the opcode for the last surviving command, together with the repeat count identifying the number of pixels that qualified before the last surviving command failed to qualify.

[0064] As an example, suppose for a current pixel, CL, CA and CO commands all qualify. The encoder records that and begins counting. In the next pixel, the encoder

determines that all still apply, and so increments the counter to two. The process continues identically until, in the seventh pixel, the CL condition no longer applies. The encoder 26 drops CL out of the running and continues incrementing the counter. Continuing, suppose in the 14th pixel, the CA condition becomes false. The CO command is the last survivor, but the encoder still does not stop counting. It continues incrementing until, suppose in the 51st pixel, the CO condition becomes false. At that point, the encoder 26 sends two consecutive bytes to the client 11: 00000001 and 00010000. The first byte indicates a CO condition (opcode=000) for what first appears to be a repeat count of three (recalling that a "zero" specifies a repeat count of two). But, when the decoder 30 looks ahead to the next byte, it sees that the consecutive CO commands are to be read together to form a ten-bit word. (Note that the decoder 30 will also look to the next byte beyond the 00010000 byte before decoding the word, to be sure that a third CO byte does not follow the second one). The ten-bit word: 0000110000 equates to a repeat count of 50. This series of two CO commands instructs the decoder to not change the next 50 pixels from the colors they were in the previously sent frame.

[0065] Once a copy command becomes the last one standing, the opcode for the next command is determined. When this last standing command no longer qualifies, the repeat count for that command is determined. At that point, the encoder also determines how many bytes are necessary to identify the repeat count. If the count can be provided in five bits, the encoder generates a one-byte command. If ten bits are required, the encoder generates a two-byte command, and so forth. This aspect of the preferred embodiment is advantageous because it capitalizes optimally on the identification of the longest possible repeat counts. In fact, one can envision other copy commands, other than CA, CL and CO, which identify pixels based on other relational aspects.

[0066] The hierarchical priorities between the CO, CL and CA commands apply if two or more of those commands are equally last standing. In that case, the encoder resorts first to the copy old command. The copy old command presents the least burden on the client because the result is only skipping over pixels. On the other hand, the client has to work to copy from above or to copy from the left. As between these two copy commands, the copy left is higher priority than copy from above, again because it

presents less of a burden to the client. With a copy left, the client only needs to read the immediately preceding pixel once and write it a number of pixels. To copy from above, however, relies on reading a number of pixels from the video line above and writing to a number of pixels.

[0067] On the other hand, if the client were implemented with hardware rather than software, the copy command priority may not matter because the hardware may be dedicated to processing commands. The preferred embodiment minimizes the load on a software client by prioritizing the copy commands.

[0068] The fourth command type (and the highest priority of non-copy commands) is the MS (make-series) command shown in Figures 3, 9 and 10. Based on analysis of the compression of typical computer screens, the make-series command ends up contributing a great deal to the efficiency of the compression. The theory on the MS command is that text, no matter what color it is in, is almost always in a two-color mode. Indeed, the inventor examined typical computer screens and determined that large portions of text and other areas of the screen can be defined with long MS commands. The MS command provides great efficiency in compressing the text portions of icons, of documents, of labels, and of tool bars. Other compression schemes either do not provide the requisite compression efficiency or do not provide the sharpness demanded by users who have to read text material on the screen.

[0069] Take, for example, the instance where a user is scrolling through text such that from one frame to the next, the text is just shifting up a little bit. From the compressor's point of view, each frame is a new group of pixels that need to be encoded. The compressor may get some repeat count efficiency by writing CO commands for areas around the text window, but when it hits the adjusted text, repeat count compression becomes inefficient because long repeat counts don't occur. The inventor has added efficiency for those text-type areas where copy commands don't work well. Exactly how those MS commands add compression efficiency will now be described.

[0070] First, like before, the three-bit opcode identifies the MS command. The first opcode bit (0) indicates that the command is not a make-pixel command. The next two bits (11) identify the command as a make-series command. Opportunities to evoke

the MS command are identified by the encoder looking ahead four pixels. The artisan should note that the copy commands require no look-ahead operation (though look-ahead operations can be added for the sake of providing additional features). With the MS command, alternatively, more or less pixels can be used for this look-ahead operation. As will be seen the number of pixels in the look ahead should be chosen strategically to be (1) large enough to ensure that repeat count coding won't be more efficient, (2) short enough to make the MS command appropriately applicable, and (3) valued as an integer that accommodates the word length being used. Solely for purposes of example herein, four pixels will be described. MS commands are invoked when the encoder determines that, within the next four pixels, two conditions occur: (1) that aCO, CL or CA command is not going to qualify, and (2) all the pixels in those next four pixels are limited to two different colors. The "extended" MS command, shown by example in Figures 9 and 10, extends the MS operation, but only the first byte includes the opcode in bits 5, 6, and 7. The extended MS command is described further below.

[0071] As previously described, the MS command is used for a series of pixels that are a combination of two different colors. The two colors that are included in the set of available colors are the color from the immediately preceding pixel (color 0) and the most recent different color pixel before that (color 1). Of course, other methods of identifying the two pixel colors for the MS command can be employed from a variety of options, including strict identification of the colors, identification from selected positions in the present frame, or the previous frame, identification from a lookup table of two-color sets, etc. In the preferred embodiment, the two colors are derived from the immediately preceding two different color pixels, which may have been encoded using make-pixel, copy-above, copy-left, or copy old commands. The MS command does not care how these two pixels got there, just that they will become the two colors for the upcoming MS command's series of pixels.

[0072] The MS command with the two-color set described above is advantageous because it does not require bytes with any color identification bits. That is, the MS commands do not include bits that identify which colors are being used, only which of the two previously identified colors are being used in the series. So, for example, when

the encoder reaches the beginning of some text, such as the top left corner of a black letter "H" on a white background, the first pixel on the top left corner of the "H" may be defined with a black MP (make-pixel) command followed by a CL (copy-left) command for a few pixels. As the top center and top right of the H are found by the encoder's look ahead, the encoder will create a make-series command because it is detecting only two colors (text and background) in the coming pixels.

[0073] As shown in Figure 9, the first MS command byte has a three-bit opcode followed by an "extension" bit (in bit position 4) that indicates that this command is extended to the next byte. If the extension bit is zero, the MS command is not extended, and it ends after the first byte. In such a case, the four "C" bits in that byte provide the two-color pattern for four pixels and then the present series ends. If, however, the extension bit is on, then another whole byte of MS data will follow the present one. Thus, in Figure 9, the second byte is an "extended command" byte. Because the extension bit is present in the preceding byte, the next byte need not include the three-bit opcode. The identity of the extended command thus comes not from an opcode in the present byte, but from the extension bit in the preceding byte. The result provides seven bits for make-series data for every byte following the first byte. Each extended command byte includes its own extension bit (in bit position 7) that identifies whether or not a next byte is an extension byte. This extension can go on as long as the E-bits are on. When the E-bit goes off, the present series stops. The series of Figure 10 indicates an example of a 13 byte long MS command that will define a series of 88 consecutive pixels.

[0074] As the decoder receives the make-series bytes, it begins immediately creating pixels for the client screen, as follows. After reading the opcode 011, the decoder realizes that a make-series is beginning. It reads the color of the preceding pixel and defines that color as "color 0". Then it reads the most recent different color pixel before that and defines that color as "color 1".

[0075] The decoder then reads the E-bit to determine whether the series is one byte, or more. Finally, the decoder reads bits 0-3, in order, and creates pixels from the two available colors based on the binary status of each pixel. For the first byte, the decoder will create four pixels. For example, if the first MS byte is 01110110 and color 0

is black and color 1 is white, the decoder will create four pixels (0110) of black, white, white, and black. Then, because the E-bit is set to 1, the decoder will look to the next byte to create seven more black and white pixels.

[0076] The first byte of an MS command, in the preferred embodiment, creates four pixels (eight bits minus three opcode bits minus one extension bit). If the encoder finds that less than four pixels are present in the series (i.e., more than two colors are present in the next four pixels), then the MS command cannot be used in the preferred embodiment. Further, if a first extension byte (a second cumulative byte) of MS command is to be used, the encoder must look ahead to find that the next seven consecutive pixels qualify for MS status (i.e., all from only two color choices and no copy command applies). Then, as shown in Figure 9, the four C-bits in the first byte will identify the first four pixels of the eleven-pixel series, and the seven C-bits in the second byte will identify the next seven pixels in the eleven-pixel series. Thereafter, new MS extension bytes are used only when whole multiples of seven pixels can be added to the series. Hence, as previously described, the encoder "looks ahead" before encoding any MS command byte, in order to: (1) determine whether the first four pixels qualify for MS treatment, and (2) determine whether additional bytes of seven pixels will qualify.

[0077] As will now be understood, the MS command defines sequential pixels, using sequential bits, such that each bit corresponds to each pixel being either color 0 or color 1. In effect, the C-bits of the MS commands are like a pixel train.

[0078] As previously described, the encoder in the MS mode is always looking ahead and won't set the E-bit unless it sees that it will have enough pixels in the coming series of pixels to fill the next seven bits of a next extension command byte. If the encoder looks ahead and encounters a color different from the two-color set, within the next seven pixels, then it ends the make-series command with the current byte (writing a stop bit into the E-bit of the current byte).

[0079] In one embodiment, the encoder is doing comparisons for all of the command types for all of the pixels all of the time. In that case, the comparisons are always running in parallel, and are always running for all commands. When one of the command types recognizes its own applicability, the encoder flags it and determines

(based on other comparisons and priorities among the commands) which of the command types is the optimum one for the present situation. In the embodiment of Figure 2, for example, the video compressor 23 is, for every single pixel, looking for the applicability of each of the five command types, and looking ahead in accordance with the MS command requirements.

[0080] The embodiments described above do not work well on the first presentation of photographs on a screen, because photographs require a relatively large number of make-pixel MP commands. Until a still photo is sent once, the encoder does not create many copy commands, which create better efficiencies. Of course, after a still photograph is initially sent to the client, the encoder will generate CO commands for those parts of the screen on subsequent frames. The present embodiments, while less applicable to photographic information, provide extraordinary efficiency in the application of computer console screens, where many vertical and horizontal lines frequently qualify for copy commands and screens include a significant amount of text.

[0081] The embodiment of the present inventions referred to as the 7-bit gray-favored color mode provides a novel and creative use of the make-pixel (MP) command vis-a-vis color and gray intensity charts. This mode aims to achieve the maximum performance from the 7-bit payload of a one-byte MP command. As shown in Figure 11, each of the incoming colors (red, green and blue) ranges in intensity anywhere from 0 (darkest) to 255 (brightest). Some existing computer console color depth reduction schemes use six total bits to define all colors (two bits are provided for red, two for blue and two for green), resulting in four different shades of red, four different shades of blue and four different shades of green. The combination of 4^3 is 64 possible color combinations.

[0082] Grays are also important in computer applications, and consist of each combination in which R, G, and B are present in equal intensity. The six-bit color scheme described above, by default, provides four possible shades of grays. While the four shades of R, G, and B may provide acceptable color depth, the limited numbers of gray shades prove unsatisfactory for gray-scale depth.

[0083] In an example embodiment (though not a limiting one), the number of colors can be increased beyond 64 while also increasing the number of gray shades by a greater proportion than the colors were increased. To do so, a "popularity of use" for all colors (including grays) is assigned based on a collection of arbitrary computer console screens, a predetermined color selection, etc. and, from that, a frequency table identifies which colors (and grays) are considered most popular. In Figure 11, the binary and decimal intensity levels (from 0-255) are shown in the left columns, followed by "popularity of use" ranking. In that column, the longer the line, the more that color was identified in the pool of typical computer screens. As shown, the zero intensity is used often, 63 and 64 are used often, 127 and 128 are used often, 191 and 193 are used often and 255 is used often.

[0084] The inventor found that grays were more popular than non-grays on typical computer screens. For example, the scroll bars are gray, the toolbars are gray, and when a "button" is pushed, the edges around buttons are changed to different shades of gray. Black and white are shades of gray and are very frequently used. Computer screens use a lot of different shades of gray, and the shade varieties are important for contrast. As color depth was taken away for video compression purposes, the first place that video quality suffered was on the grays. As it turned out, the actual colors were less critical. For example, it was less important how red a red was or how green a green was. But when depth of grays went away with the color depth reduction scheme, important contrasts like when a "button was pushed" on the screen were lost.

[0085] By looking at the popularity of colors, by providing five shades each of R, G, and B, and by finding code locations to add more grays, the present embodiment provides all of the colors needed for good color contrast while adding excellent gray scale contrast. First, a set of popular red, green, and blue intensities was selected. For the example in Figure 11, each of red, green and blue can occur in one of five of the most popular intensities: 0, 64, 128, 192 and 255. Those become the five different shades provided for each of the colors, namely five shades of red, five shades of green, and five shades of blue. The total number of colors available using those five shades is: $5^3=125$. Within that 125 colors will automatically be five shades of gray, specifically: (1) R, G,

and B all equal 0, (2) R, G, and B all equal 64, etc. Five grays is better than four, but it is still not as good as one might desire.

[0086] For that reason, additional grays can be coded into a "hidden" area of the pixel encoding. As shown in Figure 4, the MP command defines the Red, Green and Blue intensities with seven bits. 128 states (2^7) can be defined by these 7 bits but, using the five-shade popular color scheme described above, only 125 colors are identified. The present example uses the leftover three states (128 less 125) for three additional gray scales. Now, instead of five shades of gray (RGB = 0, 64, 128, 192, and 255), three additional shades of gray (RGB = 96, 160 and 224) are included. The eight grays are shown in the far right column of Figure 11.

[0087] Figure 12 is a color print of a test pattern (referred to as the 0-255 RGB+Gray test pattern) on the video screen of a computer set for 24-bit color. The test pattern has horizontal bars of pure red, pure green and pure blue, increasing from zero (darkest) to 255 (brightest). It also has a bar of pure gray (equal amounts of red, green and blue) increasing from zero to 255. Figure 13 is a color print of a resulting client screen when the "7-bit gray-favored color mode" embodiment of the present invention is employed and the source video is the test pattern shown in Figure 12. In the end, the 7-bit gray-favored color mode accurately displays the most popular five shades of red, green and blue and provides more levels of grays than the artisan would expect from 7-bits.

[0088] Compared to the prior art six-bit color schemes, the 7-bit gray-favored color mode provides better color quality, with twice as many grays (eight versus four). The 7-bit gray-favored color mode has particular application in the computer arts where high color depth is not as critical and has even more particular application in the network administration arts. Network administrators are frequently maintaining servers that are not proximate to the administrator. Still, the administrator needs to access servers and interact with the servers in real time. Getting the video from a server to a network administrator as quickly as possible after keyboard or mouse inputs is important. And prior art video schemes that return video in such poor color or gray quality, or are too slow to keep up with keyboard and mouse inputs are unacceptable. The present

compression system with the 7-bit gray-favored color mode provides excellent color quality and exceptional gray scale quality for the network administrator, who needs good video for the functional aspects of a computer interface (buttons, bars, etc.).

[0089] In another embodiment of the present invention, the color depth is dynamically increased or decreased as a function of the source video content and/or network bandwidth availability. The video compression encoder would notify the client that the length of the MP commands would be increased or decreased and all other commands would remain the same. Since the MP commands are the lowest priority and are relatively infrequent, the expansion to two or more bytes for each MP command does not dramatically increase the network traffic generated from using most computer screens. Viewing images such as photographs would increase the number of MP commands and increase the difference. Tests have shown that increasing the MP command from one to two bytes only increases traffic on typical computer screens by 30%.

[0090] In another embodiment of the present invention, network traffic can be minimized by not sending data if there are no changes to the source video from the previous frame sent. In this embodiment, when the encoder 26 recognizes that no changes have occurred, there is no need to send commands because when client 11 receives no commands, no change is made to the client screen by default. In another alternative embodiment, after some period of time (for example one minute) the server software sends a message to the client to let the client 11 know that the connection is still working and the screen has not changed.

[0091] In the embodiment described in Figure 1 and 2, the source video comes from video creation software and a video controller chip, both located in the target server 15. Another example embodiment is to have the source video come from video creation software and a video controller chip, both integrated together with the video compressor. An example of such an "embedded" fully integrated system is depicted in Figure 14.

[0092] Another alternative embodiment is to compress the video (using the same types of video commands described above) completely with software that interfaces directly with the video creation software, eliminating the need for a video controller chip.

An example of such a pure software "controller-less" embodiment is depicted in Figure 15.

[0093] In the earlier example embodiments, the command decoder was implemented with PC software. An alternative embodiment would implement the decoder completely with hardware or with a combination of hardware and a small low-cost low-performance microprocessor. This "embedded" decoder would output its video directly to a video display (without a PC or video controller chip) as shown in Figure 16. This "microclient" could also contain keyboard and mouse interface circuitry and could also be integrated into the video display. A microclient would be applicable in applications where it is desirable to have all of the workers computers out of the main work area and in a back room. In the work area, only keyboards, monitors and mice would be present on desktops. When a worker moves from one place to another, they could log onto their computer (or any other computer they were permitted to) from any microclient.

[0094] Another example aspect of the present invention will now be described with respect to Figure 17. If a second client 16 is added (the same as client 11) which also has the same client software and is also connected to the IP network, the server appliance 14 could sent the same video compression commands to both clients, permitting both clients to simultaneously "share" access to target server 15. Usually, in this "share mode," one client is accessing the server 15 and the other client is watching. The example may occur when a client 11 is employing a server and is in some operational error that the client user wishes a network administrator (who is at another location) to see. This is referred to as the "help desk" mode. The share mode is taken to greater extreme in cases where video is multicast, as to a group of trainees sitting at multiple respective client remote consoles 17 and 18.

[0095] In share mode over the Internet (especially with a large number of simultaneous users) it is advantageous to employ UDP communication instead of TCP communication. As the artisan will understand, UDP uses unacknowledged datagrams, while TCP's datagrams are acknowledged. The implosion of acknowledgements with a large number of simultaneous share mode users could flood the server appliance. The

advantage of TCP is that no data is lost because everything is sent and re-sent until acknowledged. With video, however, the user cares less about what is lost than about continuous video flow. In other words, just because the screen flickered due to a lost frame, does not mean that the user desires that the video return back to the missed frame and start over. The present invention can be employed with TCP, UDP, or any other acknowledged or unacknowledged protocol.

[0096] The applicant notes that a disadvantage of UDP protocols is that they can contribute to denial of service attacks that maliciously occur on the Internet. Because UDP is unacknowledged, traffic can flood a server with UDP datagrams. To prevent that, firewalls often block UDP. Using the present invention in the example embodiment employing UDP requires the acceptance of UDP datagrams, however training room environments and other applications for large numbers of share-mode users would typically be inside a facility behind the firewall.

[0097] In still another embodiment, data encryption is applied to the video compression commands, such that the compressed computer screens being transmitted are secure from monitoring. Any encryption technology can be employed, but an encryption technology, such as AES encryption, that could be implemented in the same video compressor 23 along with the video compression encoding would be much more desirable from an implementation viewpoint than a separate data encryption device.

[0098] The inventor presented the combination of command structures described above combined with the 7-bit gray-favored color scheme as a preferred embodiment because this combination was an optimization of trade-offs that were well suited for computer administrators working in a KVM-style server management environment. Rearranging the command opcodes and changing the color scheme can reduce network bandwidth requirements or increase color depth for other environments.

[0099] For example, if only five bits of color are used to implement the 5-bit gray-favored color mode shown in Figure 18, it would be advantageous to exchange the opcodes between the MS command and the MP command, as shown in Figure 19, because the single bit opcode would be "wasted" on an MP command with only five P bits. In that embodiment, the single bit opcode would be better used to enhance the

efficiency of the MS command. It also would eliminate the need for the MS command's extension bit (E bit), since simply sending subsequent MS commands could extend an MS command, as shown in Figures 20 and 21. This alternative combination of command structures and 5-bit color offers less color depth but improved performance on screens with a significant amount of text (because of more efficient MS commands), however it offers the same number of grays (8) as the 7-bit color mode described above.

[00100] Another embodiment optimized for applications that require more color depth, uses the same alternative arrangement of opcodes shown in Figure 19, but the MP command is either one or two bytes long as shown in Figures 22, 23 and 24. When it is two bytes long it provides 12-bit color (4 red, 4 green and 4 blue) as shown in Figure 23. When it is one byte long, it provides 4 bits of payload that define 16 shades of gray (red, green and blue all are equal) as shown in Figure 24. The "A" bit (or "all" bit) in Figure 22 indicates that all three colors are equal to the value of the "P" bits and the command is limited to one byte. Effectively, the variable length MP command is gray-favored in that less network traffic is generated from one-byte gray commands. In another embodiment, the 4-bit payload of the one-byte MP command represents the 16 most popular colors instead of 16 grays. The 16 most popular colors could be determined by recent usage statistics with MP commands, or by a pre-set list of 16 popular colors. Also, the same advantages of the more efficient MS command in the 5-bit color mode described above are included in this 12-bit color mode. The close similarity of the 5-bit and 12-bit color modes described here would allow an embodiment that dynamically switched between 5-bit and 12-bit color depending on source video content and/or available network bandwidth. Other rearrangements of the commands, similar to those shown to accommodate the 5-bit and 12-bit color modes, could also be advantageous for improved performance in other applications or other environments.

[00101] While the invention has been described in connection with what is presently considered to be the most practical and preferred embodiment, it is to be understood that the invention is not to be limited to the disclosed embodiment, but on the contrary, is intended to cover various modifications and equivalent arrangements

included within the spirit and scope of the appended claims.

CLAIMS

1. A method of compressing computer video into n-bit packets, comprising:
 - examining the color of each video pixel within selected video frames;
 - determining which of the following conditions applies for each pixel in each selected frame:
 - a first condition wherein a current pixel is unchanged versus a pixel at the same screen position in the previously selected frame;
 - a second condition wherein the current pixel is unchanged versus a pixel directly to the left of the current pixel in the current frame;
 - a third condition wherein the current pixel is unchanged versus a pixel directly above the current pixel in the current frame;
 - a fourth condition wherein a sequential series of x pixels, beginning with the current pixel, are comprised only of colors from a two-color set and x is at least (n-4); and
 - a fifth condition wherein each of the first, second, third and fourth conditions do not apply;
 - (a) if any of said first, second or third conditions apply, determining which of said conditions applies for the longest number of sequential pixels beginning with the current pixel, and
 - (i) if said determined condition is the first condition, encoding a packet including:
 - one bit indicating that said fifth condition does not apply, second and third bits together indicating that said first condition applies, and
 - (n-3) payload bits indicating the number of sequential pixels, beginning with the current pixel, that are unchanged versus pixels at the same screen positions in the previously selected frame;
 - (ii) if said determined condition is the second condition, encoding a packet including:
 - one bit indicating that said fifth condition does not apply,

second and third bits together indicating that said second condition applies, and

(n-3) payload bits indicating the number of sequential pixels, beginning with the current pixel, that can be copied from the pixel in the current frame directly to the left of the current pixel;

(iii) if said determined condition is the third condition, encoding a packet including:

one bit indicating that said fifth condition does not apply,

second and third bits together indicating that said third condition applies,

and

(n-3) payload bits indicating a number of sequential pixels, beginning with the current pixel, that can be copied from said number of pixels in the current frame directly above each of them;

(b) if each of said first, second and third conditions do not apply and the fourth condition does apply, encoding one or more n-bit packets including:

(i) a first packet, including:

one bit indicating that said fifth condition does not apply,

second and third bits together indicating that said fourth condition

applies,

a fourth bit indicating whether the first packet is followed by another packet corresponding to the said series of x pixels;

(n-4) bits indicating which colors, from said two-color set, applies to each of the first (n-4) pixels in said series of x pixels; and

(ii) if x is greater than (2n-5), one or more subsequent packets, each including:

one bit indicating whether another packet corresponding to the said series of x pixels follows; and

(n-1) bits indicating which colors, from said two-color set, applies to each of the next (n-1) pixels in said series of x pixels; and

(c) if said fifth condition applies, encoding an n-bit packet including:

one bit indicating that said fifth condition applies; and
(n-1) bits defining the color of the current pixel.

2. A method according to claim 1, further including:
determining that more than (n-1) bits are required to define pixel colors, and
wherein;
an additional packet is encoded following the packet encoded in step (c),
which include;
n bits of additional color defining bits, thus creating a total of (2n-1) bits to
define the color of the current pixel.
3. A method according to claim 2, further including:
determining that more than (2n-1) bits are required to define pixel colors,
and wherein;
one or more further packets are encoded following said additional packet
each further packet including:
n bits of additional color defining bits, thus creating a total of (Yn-1) bits
to define the color of the current pixel where Y is the sum total number of
encoded packets.
4. A method according to claim 1, wherein the encoding step (a)(i) includes
encoding packets with (n-3) payload bits defining repeat counts from two to
 $2^{(n-3)}+1$.
5. A method according to claim 4 wherein n=8 and the (n-3) payload bits define
repeat counts from two to thirty-three.
6. A method according to claim 1, further including the step of:
in encoding steps (a)(i), (a) (ii) or (a)(iii), when the required repeat count
exceeds the payload capacity of one packet, encoding y number of packets of

respective (n-3) payloads which together provide a total combined payload equal to y times (n-3) bits.

7. A method according to claim 6, wherein:

each single packet payload has the capacity to define repeat counts up to $2^{(n-3)}$; when the required repeat count exceeds $2^{(n-3)}$, a second packet is encoded to provide a $2*(n-3)$ total combined payload, and the first and second packets provide for repeat counts of $2^{(2*(n-3))}$; and

when the required repeat count exceeds both the available $2^{(n-3)}$ of a single packet and the $2^{(2*(n-3))}$ of two packets, a third packet is encoded to provide a $3*(n-3)$ total combined payload, and these first, second, and third packets provide for repeat counts up to $2^{(3*(n-3))}$.

8. A method according to claim 1, wherein the encoding step (c) includes:

defining less than $2^{(n-1)}$ colors for said current pixel with said (n-1) payload bits and defining additional gray intensities with the unused remainder of the $2^{(n-1)}$ color combinations.

9. A method according to claim 1, wherein the encoding step (c) includes using seven bits to define:

five different intensities each of red, green, and blue whereby a combination of equal intensities of red, green, and blue define five gray intensities; and

three additional gray intensities different from the five gray intensities defined by the equal intensities of red, green and blue.

10. A method according to claim 1, further comprising the step of:

decoding the packets by software resident on a personal computer.

11. A method according to claim 1, wherein the steps of examining the color of each video pixel and determining which of the conditions applies for each pixel further includes:

reading ahead of the current pixel to subsequent pixels to determine whether the fourth condition applies; and

determining whether more than one of the prioritized conditions applies for the current pixel, and if so, for the subset of applicable prioritized conditions:

(i) reading the next-subsequent pixel to determine if the subset of prioritized conditions also applies to the next-subsequent pixel;

(ii) if one or more of the subset of prioritized conditions does not apply to the next-subsequent pixel, deleting the inapplicable one or more prioritized conditions from the subset of applicable prioritized conditions; and

(iii) repeating steps (i) and (ii) until no prioritized condition still applies and thereafter:

(a) encoding the repeat count into one or more packets in accordance with the highest rated prioritized condition last left standing.

12. Apparatus for compressing computer video into n-bit packets, comprising an encoder that:

(1) examines the color of each pixel in selected video frames;

(2) determines which of the following conditions applies for the current pixel in the selected frame:

a first condition wherein a current pixel is unchanged versus a pixel at the same screen position in the previously selected frame;

a second condition wherein the current pixel is unchanged versus a pixel directly to the left of the current pixel in the current frame;

a third condition wherein the current pixel is unchanged versus a pixel directly above the current pixel in the current frame;

a fourth condition wherein a sequential series of x pixels, beginning with the current pixel, are comprised only of colors from a two-color set and x is at least (n-4); and

a fifth condition wherein each of the first, second, third and fourth conditions do not apply; and

(3) encodes a packet, which includes:

(a) if it is determined said first condition applies:

one bit indicating that said fifth condition does not apply,

second and third bits together indicating that said first condition applies, and

(n-3) payload bits indicating the number of sequential pixels, beginning with the current pixel, that are unchanged versus pixels at the same screen position in the previously selected frame;

(b) if it is determined said second condition, but not said first condition, applies:

one bit indicating that said fifth condition does not apply,

second and third bits together indicating that said second condition applies,

and

(n-3) payload bits indicating the number of sequential pixels, beginning with the current pixel, that can be copied from the pixel in the current frame directly to the left of the current pixel;

(c) if it is determined said third condition, but not said first or second conditions, applies:

one bit indicating that said fifth condition does not apply,

second and third bits together indicating that said third condition applies,

and

(n-3) payload bits indicating a number of sequential pixels, beginning with the current pixel, that can be copied from said number of pixels in the current frame directly above each of them;

(d) if it is determined said fourth condition, but not said first, second or third conditions, applies:

(i) a first packet, including:

one bit indicating that said fifth condition does not apply,

second and third bits together indicating that said fourth condition applies,

a fourth bit indicating whether the first packet is followed by another packet corresponding to the same series of x pixels; and

(n-4) bits indicating which color, from said two-color set, applies to each of the (n-4) pixels in said series of x pixels; and

(ii) if $x > (n-4)$, one or more subsequent packets, each including:

one bit indicating whether another packet corresponding to the same series of x pixels follows; and

(n-1) bits indicating which color, from said two-color set, applies to each of the (n-1) additional pixels in said series of x pixels; and

(e) if it is determined said fifth condition applies:

one bit indicating that said fifth condition applies; and

(n-1) bits defining the color for the current pixel.

13. Apparatus according to claim 12, the encoder further:

determining if more than (n-1) bits are required to define pixel colors, and wherein

one or more further packets are encoded following the packet encoded in step (e), each further packet including;

n bits of additional color defining bits, thus creating a total of $(Yn-1)$ bits to define the color of the current pixel, where Y is the sum total number of encoded packets.

14. A method of compressing video, comprising:

comparing some portions of selected video frames to other portions, and encoding said some portions based on their locational relationship to said other portions, and;

identifying still other portions of the same selected video frames as pixels comprised only of colors from a two-color set, and encoding said still other portions as a series of bits having binary states corresponding to the two colors in said two-color set.

15. A method of compressing fixed-bit RGB video, comprising:

reducing an available number of different pixel colors by:

grouping predetermined intensities of the red, green and blue components of each pixel color into intensity zones and;

adding a greater number of gray intensity zones than would be naturally created when said red, green and blue components are all in the same said intensity zones as each other.

16. Apparatus, comprising:

an encoder that creates a data packet defining a present repeated sequence of pixels by copying the color of another pixel selected based on a frame location relationship to the present repeated sequence of pixels, and selects, for the present repeated sequence of pixels, the selected frame location relationship from a group of frame location relationship types such that the selected frame location relationship will yield the longest repeated sequence of all the frame location relationship types.

17. Apparatus according to claim 16, wherein:

when said encoder finds that the length of said repeated sequence for two or more frame location relationship types are equal, selecting the selected frame location relationship based on a predetermined hierarchy between said two or more frame location relationship types.

18. Apparatus according to claim 16, wherein:

said group of frame location relationship types include:

- a pixel located to the left of said present pixel in the same frame;
- a pixel located above said present pixel in the same frame; and
- a pixel located in the previous frame at the same location as said present pixel.

19. Apparatus according to claim 16, wherein the encoder further:
makes a packet that defines color of the present pixel from a specific set of two colors.
20. Apparatus according to claim 16, wherein the encoder further:
makes a packet that defines a color of the present pixel from a test pattern of available colors.
21. A video compression routine including:
examining pixels in selected past and present video frames;
for a given said present frame and for a current pixel thereof, making fixed bit length packets that define at least the current pixel, said packets including:
(1) pixel-copy packets having at least one packet-type identifier bit and at least one repeat count identification bit;
(2) color defining packets having at least one packet-type identifier bit and at least one color identifier bit.
22. A video data compression routine according to claim 21, wherein said algorithm further includes:
for said given present frame and current pixel:
making two-color series encoded packets having at least one packet-type identifier bit and a series of binary color identifier bits corresponding to only two colors and coinciding with a color of said present pixel and a number of pixels immediately following said current pixel.
23. A video transmission system, comprising:
a video encoding routine to encode serial pixel data at a first processor according to an algorithm including, for a given set of pixels, choosing a more efficient one of:
(1) copy-pixel encoding;

- (2) individually colored pixel encoding; and
- (3) two-color series pixel encoding.

24. A video transmission system according to claim 23, wherein:
the copy-pixel encoding includes encoding based on a frame location relationship between the present pixel in the present frame and another pixel in the present frame.
25. A video transmission system according to claim 23, wherein:
the copy-pixel encoding includes encoding based on a selection of:
a relationship between the present pixel in the present frame and another pixel to the left of the present pixel in the present frame;
a relationship between the present pixel in the present frame and another pixel above the present pixel in the present frame; and
a relationship between the present pixel in the present frame and another pixel at the same location but in a previous frame.
26. A video transmission system according to claim 23, wherein:
the two-color series pixel encoding includes encoding wherein a sequential series of x pixels, beginning with the current pixel, are comprised only of colors from a two-color set.
27. A method of encoding video, comprising the steps of:
predefining a set of pixel-copy commands based on frame location relationships between the present pixel and other pixels;
for the present pixel, encoding according to a hierarchy selection from one of:
an identified repeat count associated with at least one of the pixel copy commands;
an identified series of pixels drawn from only two colors; and

an identified specific individual pixel color.

28. A method according to claim 27, wherein the encoding includes making fixed-size packets including an opcode portion and a payload portion.
29. A method according to claim 28, wherein at least some of the fixed-size packets include:
- one bit identifying the hierarchy selection associated with individually colored pixels;
 - two additional bits identifying the hierarchy selections associated with three different pixel copy commands and the two-color series pixels command; and
 - a payload of at least n-bits.
30. A method according to claim 29, wherein others of the fixed-size packets include an extension bit linking the current packet with a previous packet and including a payload of greater than n-bits.
31. A method according to claim 29, wherein others of the fixed-size packets include an extension bit linking the current packet with the next packet, which next packet then includes a payload of greater than n-bits.
32. A method of compressing a video stream to be displayed at pixel locations on a screen, comprising:
- (a) sequentially receiving frames of information from the video stream comprising sequential series of respective pixel identifiers;
 - (b) for a first received pixel identifier, comparing the first pixel identifier with;
 - (i) an adjacent-pixel identifier adjacent the first received pixel on the screen;

(ii) an above-pixel identifier above the first received pixel identifier on the screen; and

(iii) an old-pixel identifier at the same screen location as the first received pixel identifier but on a previous screen; and

(c) counting a number of pixels from the first received pixel identifier for which at least one of said conditions (i), (ii), and (iii) yields a continuous stream of identity conditions;

(d) encoding a data packet having:

a code identifying one of the three conditions (i), (ii), (iii) which provides the longest said continuous stream, and

information identifying the length of said longest continuous stream.

33. A system, comprising:

an input to receive a serial stream of information specifying a serial sequence of pixel video values in an order in which said pixel video values are to be written onto a video monitor;

an encoder continually reading the serial stream of information and determining whether a condition occurs in which a predetermined number of consecutive ones of said pixel video values consists from a group of only two color values, and when said condition occurs, encoding the predetermined number of consecutive ones of said pixel video values by creating a series command with a series of binary state values corresponding in sequence to said consecutive ones of said pixel video values and wherein one of said binary state values corresponds to one of said two pixel video values and the other of said binary state values corresponds to the other of said two pixel video values; and

said encoder encoding said serial stream of information using an alternative identification of said pixel video values when said encoder determines that said condition does not occur.

34. A system according to claim 33, wherein the encoder creates an instruction regarding the identity of said two color values followed by the series of binary states corresponding in sequence to said consecutive ones of said pixel video values.
35. A system according to claim 34, wherein:
said instruction regarding the identity is an instruction making reference to prior pixel video values in said serial stream of information.
36. A system according to claim 33, wherein:
the encoder corresponds with a decoder and said encoder provides no instruction to the decoder regarding the identity of the two color values:
when said condition occurs, the encoder provides an instruction identifying that the type of command being communicated is a command corresponding to a two color series; and
said decoder identifies the two color values based on a pre-established prior location of the two color values in the serial stream of information.
37. A system according to claim 33, wherein:
the alternative identification of said pixel video values includes a series of corresponding command packets digitally identifying consecutive corresponding ones of said pixel video values.
38. A system, comprising:
an input to receive a serial stream of information specifying a serial sequence of pixel video values in an order in which said pixel video values are to be written onto a video monitor;
an encoder continually reading the serial stream of information and determining whether a condition occurs in which a current pixel video value in the serial sequence of pixel video values is identical to a prior pixel video value

located above said current pixel video value relative to their display on a computer monitor, and when said condition occurs, issuing a command to copy the prior pixel video value to determine the current pixel video value; and

said encoder encoding said serial stream of information using an alternative identification of said pixel video values when said encoder determines that said condition does not occur.

39. A system according to claim 38, wherein:

the alternative identification of said pixel video values includes a series of corresponding command packets digitally identifying consecutive corresponding ones of said pixel video values.

40. A method of transmitting frames of video information, comprising the steps of:

receiving serial video information;

creating commands encoding the serial video information into a series of video command functions by analyzing change relationships between current values of said serial video information and prior values of said serial video information; and

dynamically and automatically adjusting a frame transmission rate in accordance with degrees of said change relationships.

41. A method according to claim 40, wherein said prior values are encoded to optimize encoding by selecting for given pixels the more efficient of (1) commands to copy prior pixel values and (2) commands to identify a current pixel value by color definition.

42. A method according to claim 40, wherein the step of dynamically and automatically adjusting the frame transmission rate in accordance with degrees of said change relationships includes, when an available transmission capacity

reduces, steadily maintaining a resolution quality of said video information while adjusting said frame transmission rate to accommodate said reduction in available capacity.

43. A method of encoding video information by:

providing color packets with 5 bits yielding 32 words for color identification;

predetermining 3 red values, 3 blue values and 3 green values for a total of 27 possible color combinations;

correlating 32 words with the 27 possible color combinations yielding 5 remaining words;

predetermining 5 extra gray values; and

correlating the remaining 5 words with the 5 extra gray values.

44. A method according to claim 43, further comprising:

providing color packets with greater than 5 bits; and

dynamically selecting the 5 bit packets for lower color resolutions and the higher bit packets for higher color resolutions.

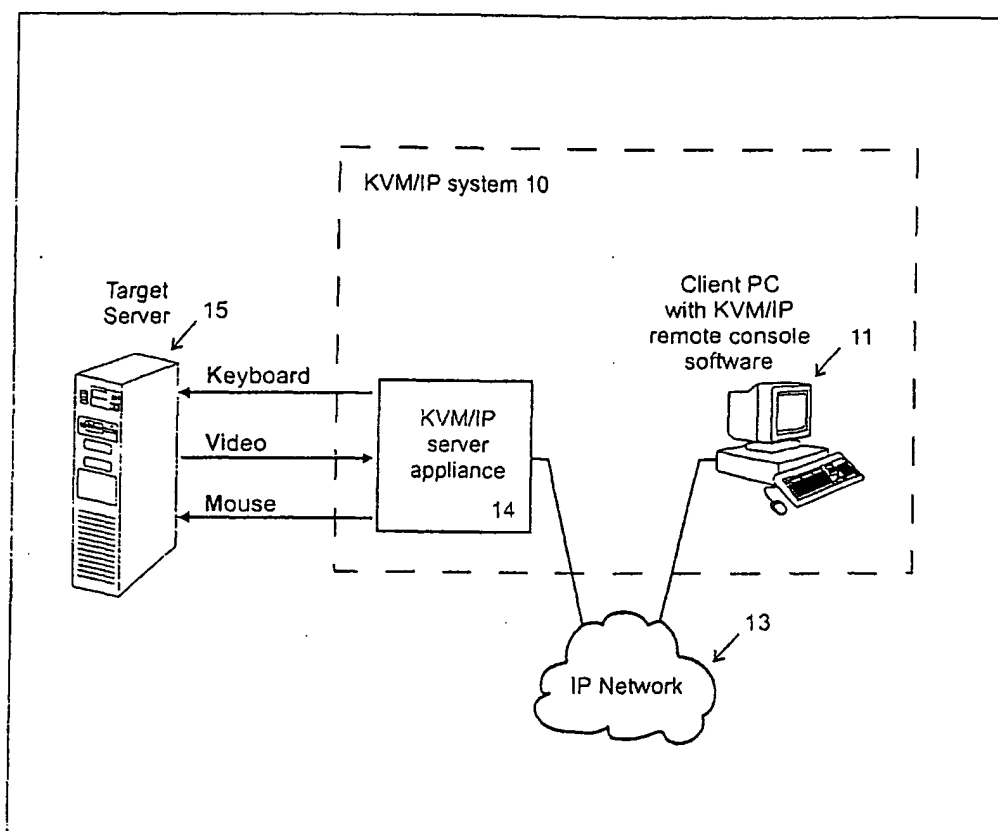


Figure 1.

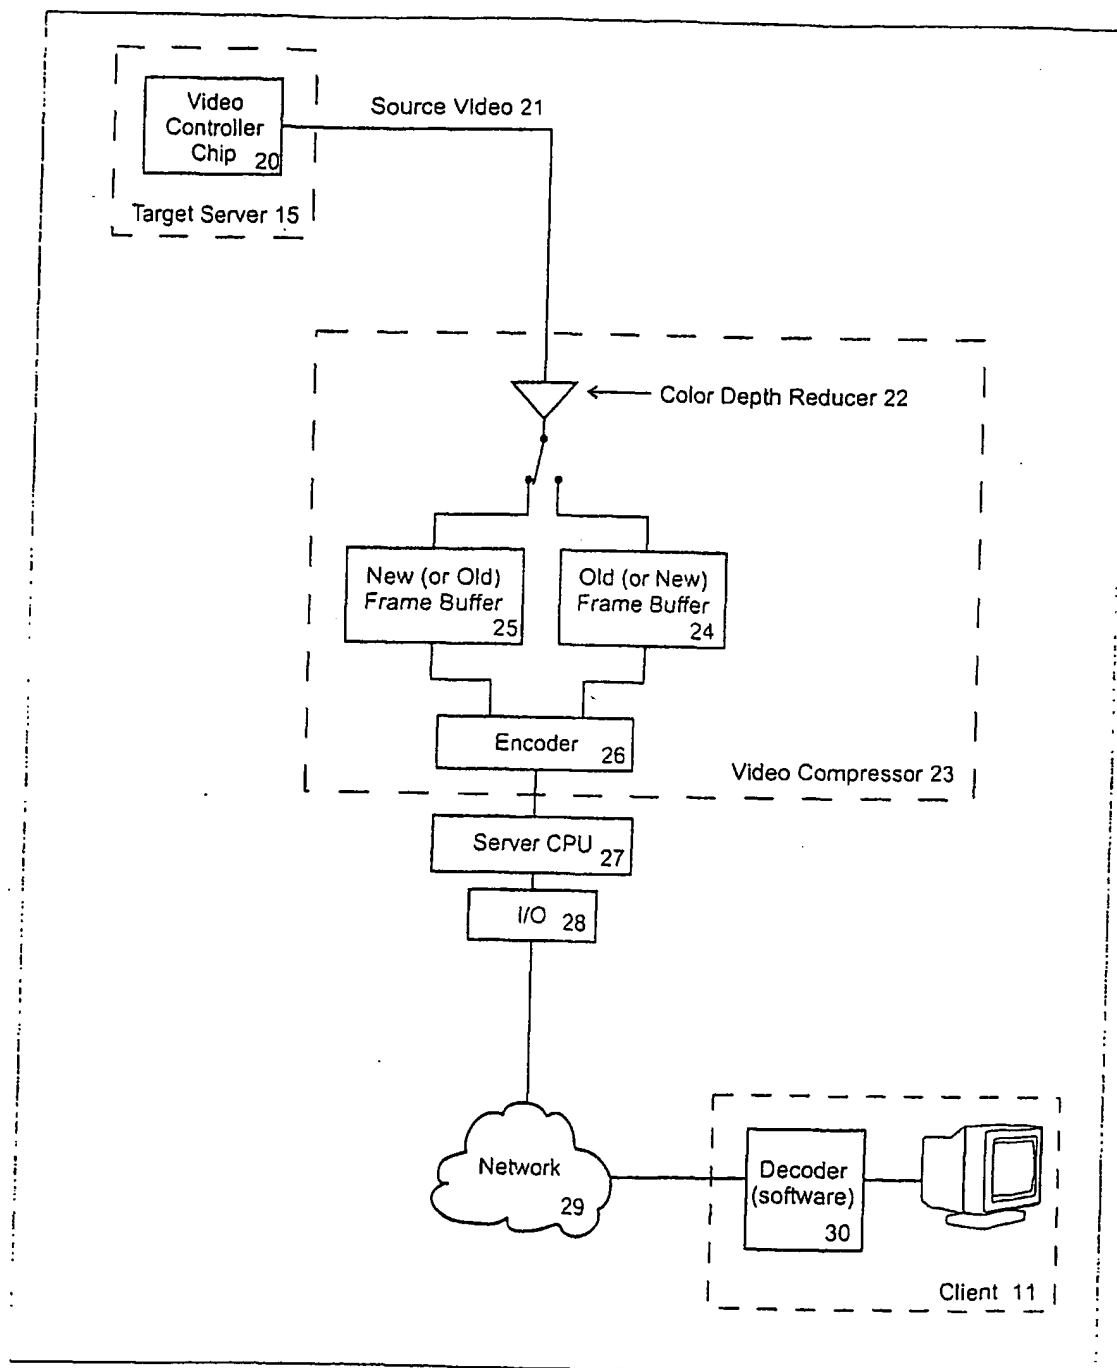


Figure 2.

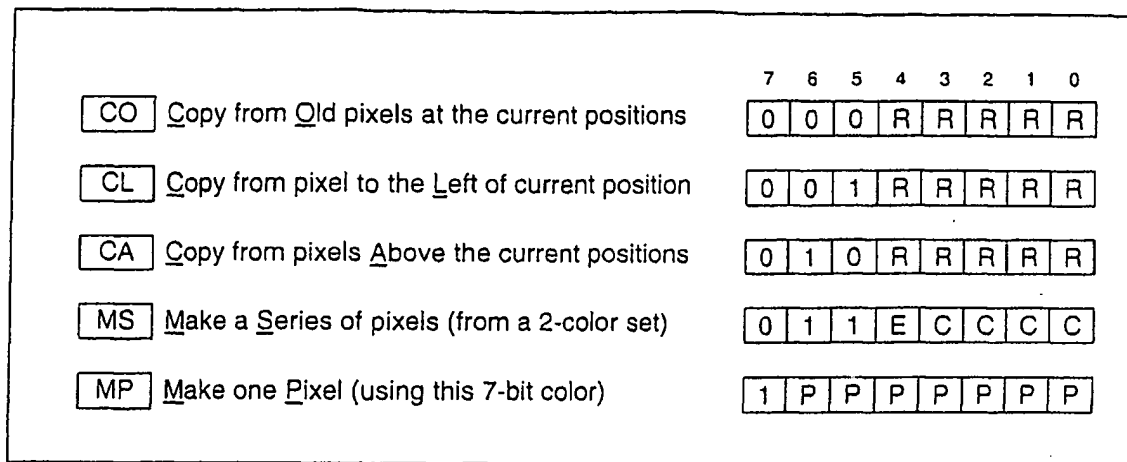


Figure 3




Figure 4




Figure 5




Figure 6




Figure 7




Figure 8




Figure 9

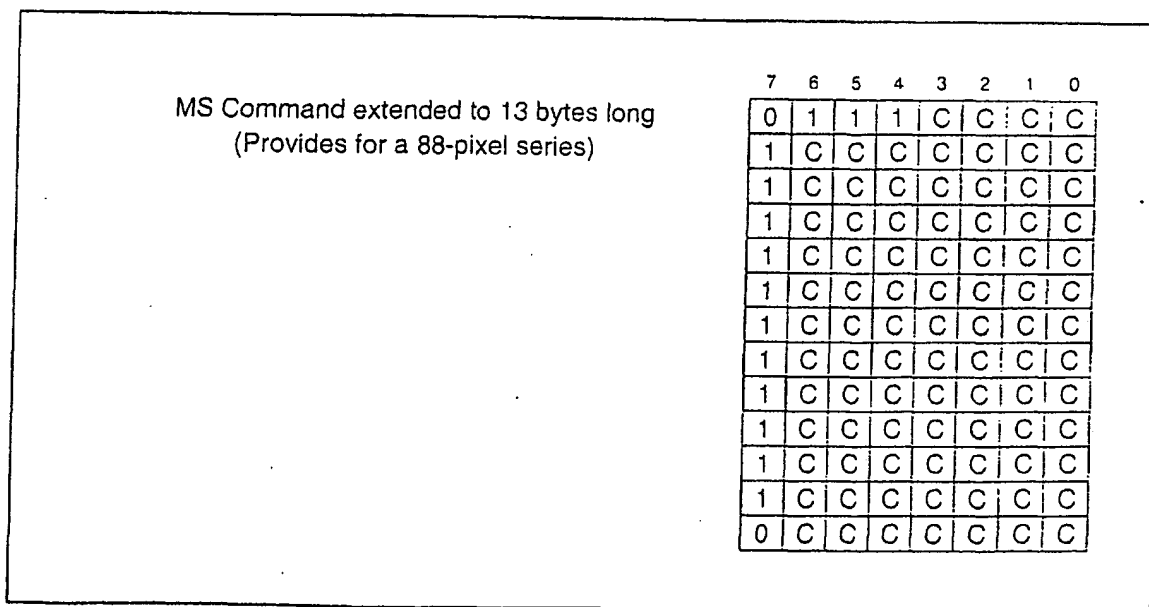


Figure 10

7-bit Gray-favored Color Mode

Intensity of source video RGB signals				Zones for all non-gray colors		Zones for all grays	
Binary	Decimal	Popularity of use		Source range	Client screen	Source range	Client screen
7 6 5 4 3		Less used	More used				
1 1 1 1 1	248 - 255		255	255	255	255	255
1 1 1 1 0	240 - 247	240					
1 1 1 0 1	232 - 239	239				232	
1 1 1 0 0	224 - 231	224				231	224
1 1 0 1 1	216 - 223	223		216		216	
1 1 0 1 0	208 - 215	208		215		215	
1 1 0 0 1	200 - 207	207					
1 1 0 0 0	192 - 199		192		192		192
1 0 1 1 1	184 - 191		191			184	
1 0 1 1 0	176 - 183	176				183	
1 0 1 0 1	168 - 175	175					
1 0 1 0 0	160 - 167	160					160
1 0 0 1 1	152 - 159	159		152		152	
1 0 0 1 0	144 - 151	144		151		151	
1 0 0 0 1	136 - 143	143					
1 0 0 0 0	128 - 135		128		128		128
0 1 1 1 1	120 - 127		127			120	
0 1 1 1 0	112 - 119	112				119	
0 1 1 0 1	104 - 111	111					
0 1 1 0 0	96 - 103	96					96
0 1 0 1 1	88 - 95	95		88		88	
0 1 0 1 0	80 - 87	80		87		87	
0 1 0 0 1	72 - 79	79					
0 1 0 0 0	64 - 71		64		64		64
0 0 1 1 1	56 - 63		63				
0 0 1 1 0	48 - 55	48					
0 0 1 0 1	40 - 47	47					
0 0 1 0 0	32 - 39		32				
0 0 0 1 1	24 - 31		31	24		24	
0 0 0 1 0	16 - 23	16		23		23	
0 0 0 0 1	8 - 15	15					
0 0 0 0 0	0 - 7		0	0	0	0	0

Figure 11




Figure 12




Figure 13

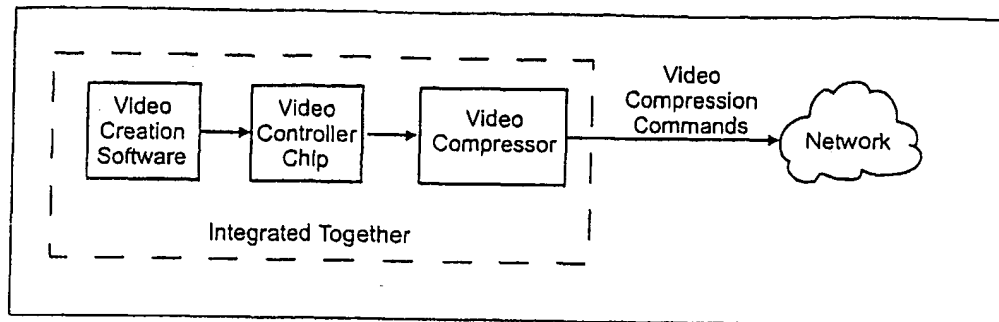


Figure 14.

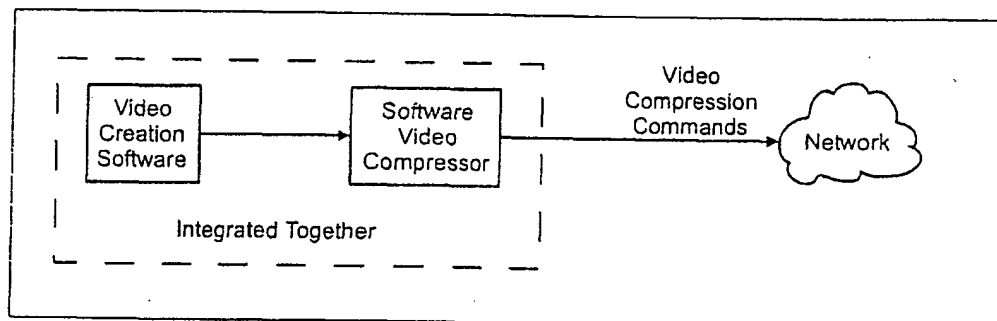


Figure 15.

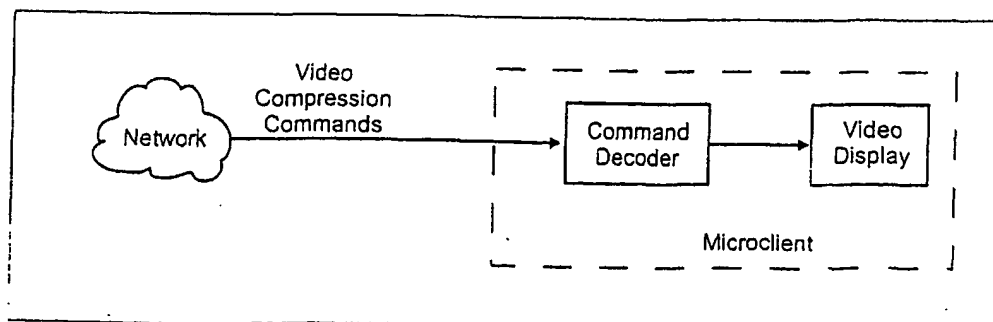


Figure 16.

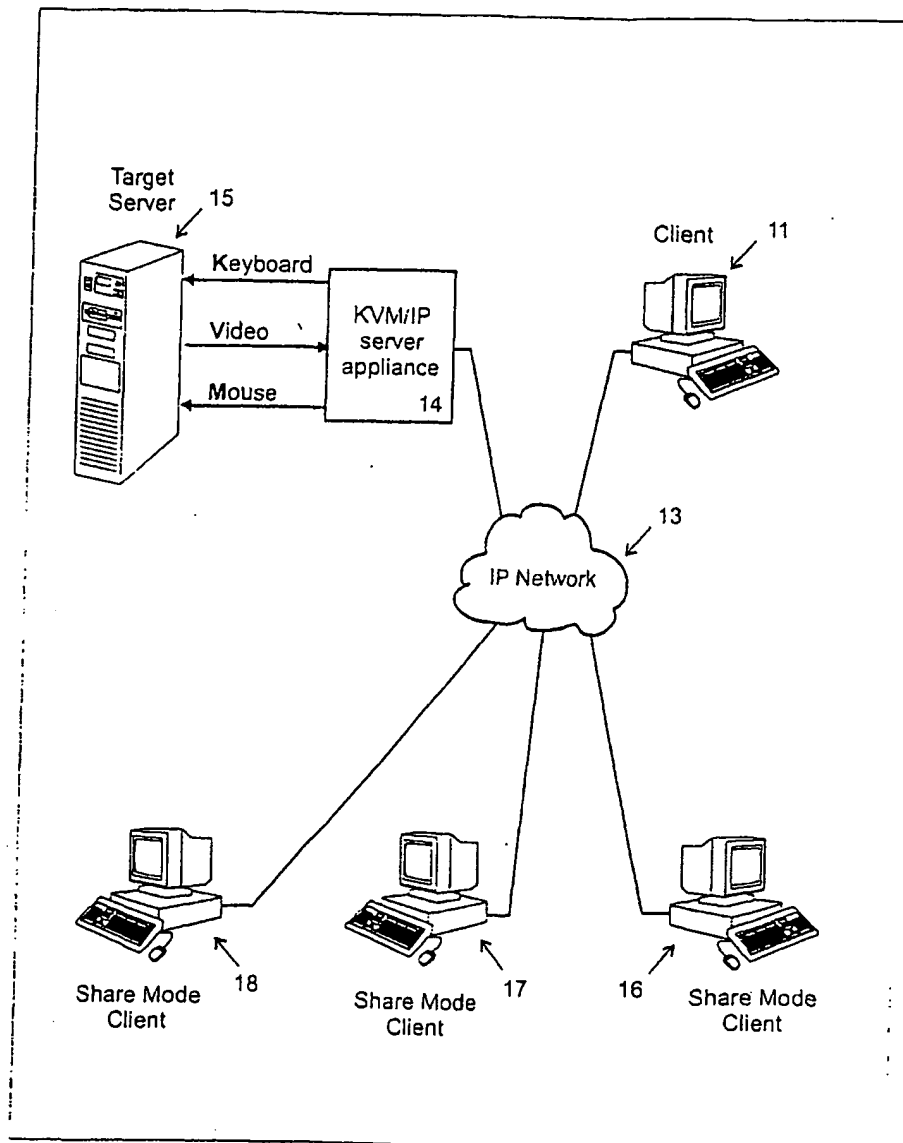


Figure 17.

5-bit Gray-favored Color Mode

Intensity of source video RGB signals				Zones for all non-gray colors		Zones for all grays	
Binary	Decimal	Popularity of use		Source range	Client screen	Source range	Client screen
7 6 5 4 3		Less used	More used				
1 1 1 1 1	248 - 255		255	255	255	255	255
1 1 1 1 0	240 - 247	240					
1 1 1 0 1	232 - 239	239				232	
1 1 1 0 0	224 - 231		224			231	224
1 1 0 1 1	216 - 223		223				
1 1 0 1 0	208 - 215	208					
1 1 0 0 1	200 - 207	207				200	
1 1 0 0 0	192 - 199		192			199	192
1 0 1 1 1	184 - 191		191				
1 0 1 1 0	176 - 183	176					
1 0 1 0 1	168 - 175	175		168		168	
1 0 1 0 0	160 - 167		160	167		167	160
1 0 0 1 1	152 - 159		159				
1 0 0 1 0	144 - 151	144					
1 0 0 0 1	136 - 143	143				136	
1 0 0 0 0	128 - 135		128		128	135	128
0 1 1 1 1	120 - 127		127				
0 1 1 1 0	112 - 119	112					
0 1 1 0 1	104 - 111	111				104	
0 1 1 0 0	96 - 103		96			103	96
0 1 0 1 1	88 - 95		95				
0 1 0 1 0	80 - 87	80					
0 1 0 0 1	72 - 79	79				72	
0 1 0 0 0	64 - 71		64			71	64
0 0 1 1 1	56 - 63		63				
0 0 1 1 0	48 - 55	48					
0 0 1 0 1	40 - 47	47		40		40	
0 0 1 0 0	32 - 39		32	39		39	
0 0 0 1 1	24 - 31		31				
0 0 0 1 0	16 - 23	16					
0 0 0 0 1	8 - 15	15					
0 0 0 0 0	0 - 7		0	0	0	0	0

Figure 18

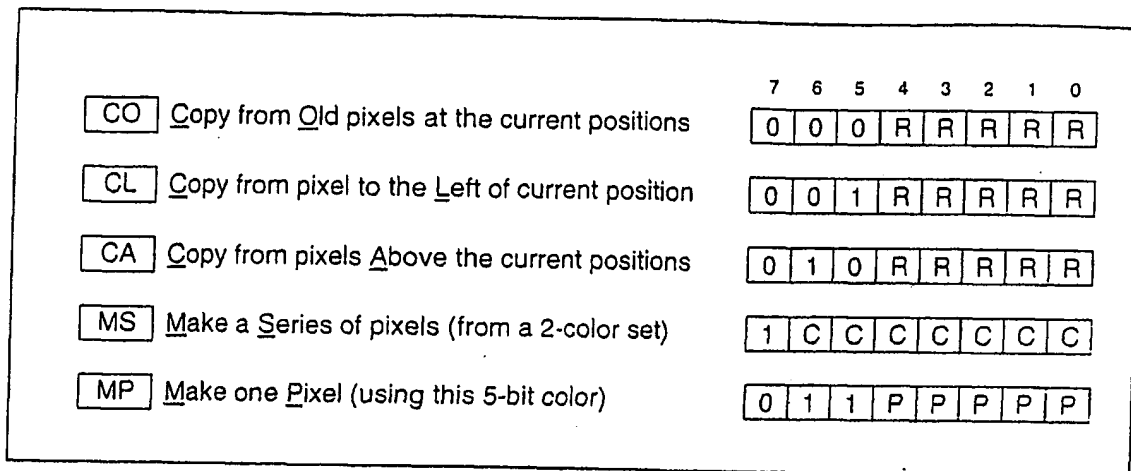


Figure 19




Figure 20




Figure 21




Figure 22




Figure 23




Figure 24

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US03/10488

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : H04B 1/66

US CL : 375/240.01

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 375/240.01

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim N
A	US 6,233,226 B1 (GRINGERI et al) 15 May 2001, fig. 1, col. 5-23	1-44
A	US 5,519,874 A (YAMAGISHI et al) 21 May 1996, fig. 1, col. 2-13	1-44
A	US 5,526,024 A (GAGLIANELLO et al) 11 June 1996, col. 2-7	1-44

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubt on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"G" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

15 JUNE 2003

Date of mailing of the international search report

28 JUL 2003

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

CHRIS KELLEY

Telephone No. (703) 305-4500